

A 594 Gbps LDPC Decoder Based on Finite-Alphabet Message Passing

Reza Ghanaatian, Alexios Balatsoukas-Stimming, Christoph Müller, Michael Meidlinger, Gerald Matz, Adam Teman, and Andreas Burg

Abstract—An ultra-high throughput low-density parity check (LDPC) decoder with an unrolled full-parallel architecture is proposed, which achieves the highest decoding throughput compared to previously reported LDPC decoders in literature. The decoder benefits from a serial message-transfer approach between the decoding stages to alleviate the well-known routing congestion problem in parallel LDPC decoders. Furthermore, a finite-alphabet message passing algorithm is employed to replace the variable node update rule of standard min-sum decoders with optimized look-up tables, designed in a way that maximize the mutual information between decoding messages. The proposed algorithm results in an architecture with reduced bit-width messages, leading to a significantly higher decoding throughput and to a lower area as compared to a min-sum decoder when serial message-transfer is used. The architecture is placed and routed for the standard min-sum reference decoder and for the proposed finite-alphabet decoder using a custom pseudo-hierarchical backend design strategy to further alleviate routing congestions and to handle the large design. Post-layout results show that the finite-alphabet decoder with the serial message-transfer architecture achieves a throughput as large as 594 Gbps with an area of 16.2 mm² and dissipates an average power of 22.7 pJ per decoded bit in a 28 nm FD-SOI library. Compared to the reference min-sum decoder, this corresponds to 3.3 times smaller area and 2 times better energy efficiency.

Index Terms—Low-density parity check code, min-sum decoding, unrolled architecture, finite-alphabet decoder, 28 nm FD-SOI.

I. INTRODUCTION

LOW DENSITY PARITY CHECK (LDPC) codes have become the coding scheme of choice in high data-rate communication systems after their re-discovery in the 1990s [1], due to their excellent error correcting performance along with the availability of efficient high-throughput hardware implementations in modern CMOS technologies. LDPC codes are commonly decoded using iterative message passing (MP) algorithms in which the initial estimations of the bits are improved by a continuous exchange of messages between decoder computation nodes. Among the various MP decoding algorithms, the *min-sum* (MS) decoding algorithm [2] and its variants (e.g., offset MS, scaled MS) are the most common choices for hardware implementation. LDPC decoder hardware implementations traditionally start from one of these established algorithms (e.g., MS decoding), where the

exchanged messages represent log likelihood ratios (LLRs). These LLRs are encoded as fixed point numbers in two's-complement or sign-magnitude representation, using a small number of uniform quantization levels, in order to realize the message update rules with low-complexity conventional arithmetic operations.

Recently, there has been significant interest in the design of *finite-alphabet* decoders for LDPC codes [3]–[7]. The main idea behind finite-alphabet LDPC decoders is to start from one or multiple arbitrary message alphabets, which can be encoded with a bit-width that is acceptable from an implementation complexity perspective. The message update rules are then crafted as generic mapping functions to operate on this alphabet. The main advantage of such finite-alphabet decoders is that the message bit-width can be reduced significantly with respect to a conventional decoder, while maintaining the same error-correcting performance [5], [6]. The downside of this approach is that the message update rules of finite-alphabet decoders can usually not be described using fast and area-efficient standard arithmetic circuits.

Different hardware architectures for LDPC decoders have been proposed in the literature in order to fulfill the power and throughput requirements of various standards. More specifically, varying degrees of resource sharing result in flexible decoders with low area requirements. On the one hand, *partial-parallel* LDPC decoders [8], [9] and *block-parallel* LDPC decoders [10], [11] are designed for medium throughput, with modest silicon area. *Full-parallel* [12], [13] and *unrolled* LDPC decoders [5], [14], on the other hand, achieve very high throughput (in the order of several tens or hundreds of Gbps) at the expense of large area requirements. Even though, in principle, LDPC decoders are massively parallelizable, the implementation of ultra-high speed LDPC decoders still remains a challenge, especially for long LDPC codes with large node degrees [15], mainly due to severe routing problems.

Contributions: In this paper, we propose an *unrolled full-parallel* architecture based on serial transfer of the messages to enable an ultra-high throughput implementation of LDPC decoders for codes with large node degrees. Moreover, we employ a finite-alphabet LDPC decoding algorithm in order to decrease the required quantization bit-width, and thus, to increase the throughput, which is limited by the serial message transfer. We also adopt an efficient pseudo-hierarchical flow for the physical implementation of the proposed decoder. To the best of our knowledge, by combining the aforementioned techniques, we present the fastest fully placed and routed LDPC decoder in the literature.

R. Ghanaatian, A. Balatsoukas-Stimming, C. Müller, A. Teman, and A. Burg are with the Telecommunications Circuits Laboratory, EPFL, Lausanne, Switzerland (email: {reza.ghanaatian, alexios.balatsoukas, christoph.mueller, adam.teman, andreas.burg}@epfl.ch).

M. Meidlinger and G. Matz are with the Vienna University of Technology, Vienna, Austria (email: {mmeidlin, gmatz}@nt.tuwien.ac.at).

Outline: The remainder of this paper is organized as follows: Section II gives an introduction to decoding of LDPC codes, as well as more details on existing high-throughput implementations of LDPC decoders. Section III describes our proposed ultra-high throughput decoder architecture that employs a serial message-transfer technique. In Section IV, our algorithm to design a finite-alphabet decoder with non-uniform quantization is explained and applied to the serial message-transfer decoder of Section III. Section V describes our proposed approach for the physical implementation and the timing and area optimization of serial message-transfer decoders. Finally, Section VI analyzes the implementation results, and Section VII concludes the paper.

II. BACKGROUND

A. LDPC Codes and Decoding Algorithms

A binary LDPC code is a set of codewords which are defined through an $M \times N$ binary-valued sparse parity check matrix as:

$$\{\mathbf{c} \in \{0, 1\}^N \mid \mathbf{H}\mathbf{c} = \mathbf{0}\}, \quad (1)$$

where all operations are performed modulo 2. If the parity check matrix contains exactly d_v ones per column and exactly d_c ones per row, the code is called a (d_v, d_c) -regular LDPC code. Such codes are usually represented with a *Tanner graph*, which contains N variable nodes (VNs) and M check nodes (CNs) and VN n is connected to CN m if and only if $\mathbf{H}_{mn} = 1$.

LDPC codes are traditionally decoded using MP algorithms, where information is exchanged as messages between the VNs and the CNs over the course of several decoding iterations. At each iteration the message from VN n to CN m is computed using a mapping $\Phi_v : \mathbb{R}^{d_v} \rightarrow \mathbb{R}$, which is defined as:

$$\mu_{n \rightarrow m} = \Phi_v(L_n, \bar{\mu}_{\mathcal{N}(n) \setminus m \rightarrow n}), \quad (2)$$

where $\mathcal{N}(n)$ denotes the neighbors of node n in the Tanner graph, $\bar{\mu}_{\mathcal{N}(n) \setminus m \rightarrow n} \in \mathbb{R}^{d_v-1}$ is a vector that contains the incoming messages from all neighboring CNs except m , and $L_n \in \mathbb{R}$ denotes the channel LLR corresponding to VN n . Similarly, the CN-to-VN messages are computed using a mapping $\Phi_c : \mathbb{R}^{d_c-1} \rightarrow \mathbb{R}$, which is defined as:

$$\bar{\mu}_{m \rightarrow n} = \Phi_c(\mu_{\mathcal{N}(m) \setminus n \rightarrow m}). \quad (3)$$

Fig. 1 illustrates the message updates in the Tanner graph. In addition to Φ_v and Φ_c , a third mapping $\Phi_d : \mathbb{R}^{d_v+1} \rightarrow \{0, 1\}$ is needed to provide an estimate of the transmitted codeword bits in the last VN iteration based on the incoming CN messages and the channel LLR L_n according to:

$$\hat{c}_n = \Phi_d(L_n, \bar{\mu}_{\mathcal{N}(n) \rightarrow n}). \quad (4)$$

Messages are exchanged until a valid codeword has been decoded or until the maximum number of iterations I has been reached.

For the widely used MS algorithm, the mappings (2) and (3) are:

$$\Phi_v^{\text{MS}}(L, \bar{\mu}) = L + \sum_i \bar{\mu}_i, \quad (5)$$




Fig. 1: (a) VN update and (b) CN update for $\mathcal{N}(n) = \{m, m_1, \dots, m_{d_v-1}\}$ and $\mathcal{N}(m) = \{n, n_1, \dots, n_{d_c-1}\}$.

and,

$$\Phi_c^{\text{MS}}(\mu) = \text{sign } \mu \min |\mu|, \quad (6)$$

where $\min |\mu|$ denotes the minimum of the absolute values of the vector components and $\text{sign } \mu = \prod_j \text{sign } \mu_j$. The decision mapping Φ_d is defined as:

$$\Phi_d^{\text{MS}}(L, \bar{\mu}) = \frac{1}{2} \left(1 - \text{sign} \left(L + \sum_i \bar{\mu}_i \right) \right). \quad (7)$$

B. High Throughput LDPC Decoders

Several high throughput LDPC decoders have been developed during the past decade in order to satisfy the high data rate requirements of optical and high-speed Ethernet networks. These decoders usually rely on a full-parallel isomorphic [16] architecture and a flooding schedule, which directly maps the algorithm for *one* iteration to hardware. More specifically, the CN and VN update equations are directly mapped to M CN and N VN processing units and a hard-wired routing network is responsible for passing the messages between them. From an implementation perspective, while such an architecture enables a very high throughput by fully exploiting the inherent parallelism of each iteration, the complexity of the highly unstructured routing network turns out to be a severe bottleneck. In addition to this routing problem, such full-parallel decoders usually require one or two clock cycles for each iteration and in the worst case as many cycles as the maximum number of iterations for each codeword, which is another throughput limitation factor.

Several solutions have been proposed to alleviate the routing problem in full-parallel decoders on both an architectural and an algorithmic level. The authors of [17], [18] suggest using a bit-serial architecture, which only requires a single wire for each variable-to-check and check-to-variable connection. While this approach can reduce the routing congestion, it also leads to a significant reduction in the decoding throughput. The decoder in [18], for example, only achieves a throughput of 3.3 Gbps when implemented using a 130 nm CMOS technology. Another architectural technique is reported in [13], where

the long wires of the decoder are partitioned into several short wires with pipeline registers. As a result, the critical path is broken down into shorter paths, but the decoder throughput is also affected since more cycles are required to accomplish each iteration. Nevertheless, the decoder of [13] is still able to achieve 13.2 Gbps in 90 nm CMOS with 16 iterations.

On an algorithmic level, the authors of [19] propose an MP algorithm, called *MS split-row threshold*, which uses a column-wise division of the \mathbf{H} matrix into S_{pn} partitions. Each partition contains N/S_{pn} VNs and M CNs, and global interconnects are minimized by only sharing the minimum signs between the CNs of each partition. This algorithmic modification was used to implement a full-parallel decoder for the challenging (2048, 1723) LDPC code in 65 nm CMOS, which achieves a throughput of 36.2 Gbps with 11 iterations. Another decoder, reported in [20], uses a hybrid hard/soft decoding algorithm, called *differential binary MP* algorithm, which reduces the interconnect complexity at the cost of some error-correcting performance degradation. A full-parallel (2048, 1723) LDPC decoder using this algorithm was implemented in 65 nm CMOS technology, achieving up to 126 Gbps [21].

To solve the problem of throughput limitations in full-parallel decoders from potentially using multiple iterations for decoding, the work of [14] presents an unrolled full-parallel LDPC decoder. In the proposed architecture, each decoding iteration is mapped to distinct hardware resources, leading to a decoder with I iterations that can decode one codeword per clock cycle, at the cost of significantly increased area requirements with respect to non-unrolled full-parallel decoders. This unrolled architecture achieves a throughput of 161 Gbps for a (672, 546) LDPC code with $d_v = 3$ and $d_c = 6$, when implemented in a 65 nm CMOS technology. It is notable that an unrolled decoder has 50% reduced wires between adjacent stages compared to a non-unrolled decoder since one stage of variable nodes is connected to one stage of check nodes with uni-directional data flow per decoding iteration. Even though this measure leads to a lower routing congestion, it is still challenging to fully place and route such a decoder. This routing issue becomes more and more severe when considering longer LDPC codes and especially with increasing CN and VN degrees to achieve better error-correcting performance, as required in wireline applications such as for the (2048, 1723) code with $d_v = 6$ and $d_c = 32$ used in the IEEE 802.3an standard [15].

III. SERIAL MESSAGE-TRANSFER LDPC DECODER

Unrolled full-parallel LDPC decoders provide the ultimate throughput with smaller routing congestion than full-parallel decoders. However, they are still not trivial to implement for long LDPC codes with high CN and VN degrees, which suffer from severe routing congestion. Hence, in this section, we propose an unrolled full-parallel LDPC decoder architecture that employs a serial message-transfer technique between CNs and VNs. This architecture is similar to the bit-serial implementations of [17], [18] in the way the messages are transferred; however, as we shall see later, it differs in the fact

that it is unrolled and in the way the messages are processed in the CNs and VNs.

A. Decoder Architecture Overview

An overview of the proposed unrolled serial message-transfer LDPC decoder architecture is shown in Fig. 2. As with all unrolled LDPC decoders, each decoding iteration is mapped to a distinct set of N VN and M CN units, which form a processing pipeline. In essence, the unrolled LDPC decoder is a systolic array, in which a new set of N channel LLRs is read in every clock cycle and a decoded codeword is put out every clock cycle.

Even though both the CNs and VNs can compute their outgoing messages in a single clock cycle, in the proposed serial message-transfer architecture each message is transferred one bit at a time between the consecutive variable node and check node stages over the course of Q_{msg} clock cycles, where Q_{msg} is the number of bits used for the messages. More specifically, each CN and VN unit contains a serial-to-parallel (S/P) and parallel-to-serial (P/S) conversion unit at the input and output, respectively, which are clocked Q_{msg} times faster than the processing clock to collect and transfer messages serially, while keeping the overall decoding throughput constant. More details on the architecture of the CN and VN units as well as our serial message-transfer mechanism are provided in the sequel.

B. Decoder Stages

The unrolled LDPC decoder, illustrated in Fig. 2, consists of three types of processing stages, which are described in more detail below. We note that the CN/VN processors of this reference decoder are similar to those of a standard MS decoder, and our modifications for these parts are discussed in Section IV.

1) *Check Node Stage*: Each check node stage consists of M CN units, each of which contains three components: a CN processor, d_c S/P units for the d_c input messages, and d_c P/S units for the d_c output messages. Moreover, the complete check node stage contains a register bank that is used to store the channel LLRs, which are not directly needed by the check node stage, but nevertheless must be forwarded to the following variable node stage, and thus need to be buffered. Hence, no S/P and P/S units are required for the channel LLR buffers in the check node stages as they are still read and put out serially.

The main task of each CN processor is to identify the two smallest absolute values among all d_c input messages that are used as the minima in order to efficiently calculate (6) for all outputs. Therefore, the CN processor is essentially composed of a pruned sorting tree and an output unit that selects the first or the second minimum, along with the appropriate sign for each output. The sorting unit consists of 4-input and 2-output compare-and-select (CAS) units that form a pruned sorter (tree) of depth $\log_2(\frac{d_c}{2})$.

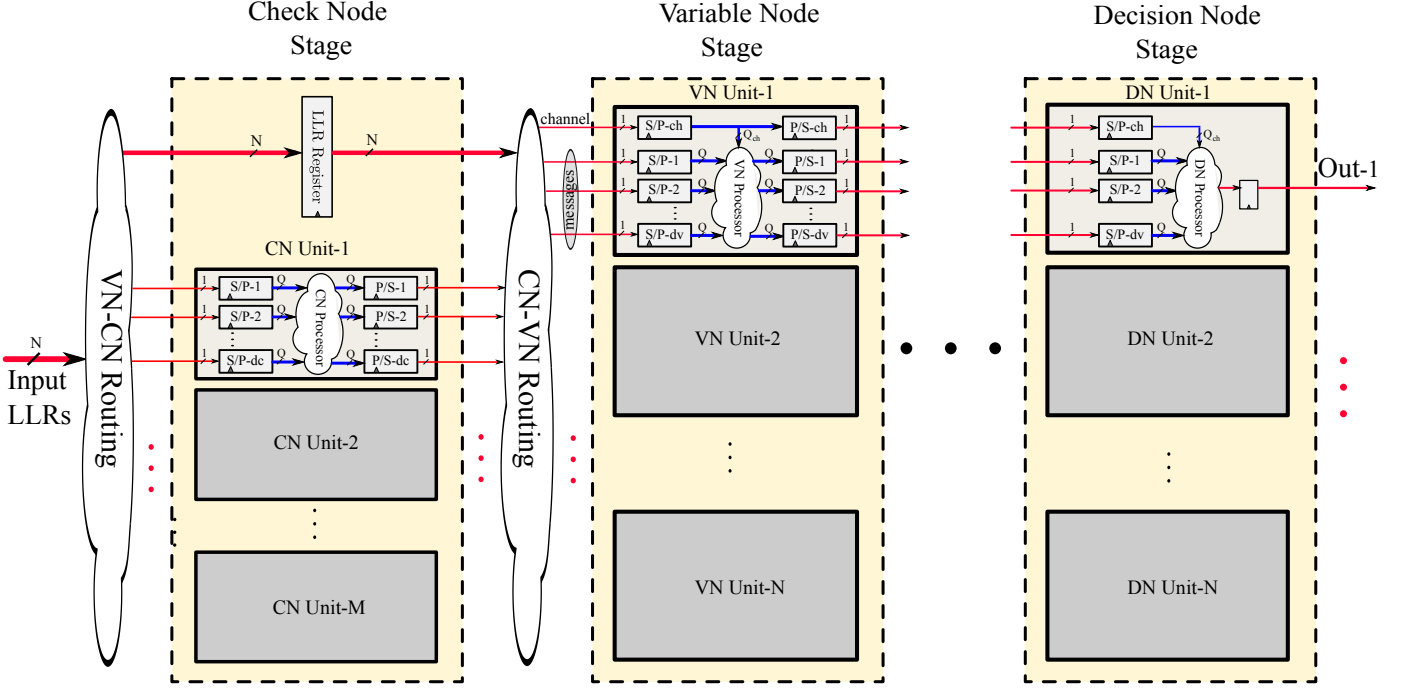


Fig. 2: Serial message-transfer decoder architecture.

2) *Variable Node Stage:* Each variable node stage consists of N VN units, each of which contains three components: a VN processor, d_v S/P units for the d_v input messages, and d_v P/S units for the d_v output messages. Moreover, the variable node stage contains N S/P and N P/S units for the channel LLRs.

The VN processor implements the update rule (5) by adding up all d_v incoming LLR messages and the channel LLR using an adder tree. An output unit subtracts each input from the aforementioned sum in order to efficiently compute all output messages.

3) *Decision Node Stage:* The last variable node stage is called a decision node stage because it is responsible for taking the final hard decisions on the decoded codeword bits. The structure of this stage is similar to a variable node stage, but a decision node (DN) is a simpler version of the VN processor in the sense that after computing the sum of all inputs, its sign is output as the decision bit and no subtraction is required. Moreover, no P/S units are required since the decision node stage is the last stage in the decoder pipeline and the output is only one bit per DN unit.

C. Message Transfer Mechanism

One of the modifications, compared to [14], is the serial transfer of the channel and message LLRs between the stages of the decoder, which reduces the required routing resources by factor of Q_{msg} . This modification is applied to make the placement and routing of the decoder more feasible, especially for large values of d_v and d_c .¹ To this end, as explained in the previous section, a S/P and a P/S shift register are added to each input and each output of the CN and VN units,

as illustrated in Fig. 3. We see that the S/P unit consists of a $(Q_{\text{msg}}-1)$ -bit shift register and Q_{msg} memory registers, while the P/S unit has Q_{msg} registers with multiplexed inputs. The serial messages are transferred with a clock, denoted by CLK_F , that is Q_{msg} times faster than the slow processing clock, denoted by CLK_S . More specifically, at each CN unit and VN unit input, data is loaded serially into S/P shift register using the fast CLK_F , and after the Q_{msg} -th cycle all message bits are stored in memory registers, clocked by the slow CLK_S . The CN/VN processing can then be performed in one CLK_S cycle and the output messages are saved in the output P/S shift registers and transferred serially to the next stage using again CLK_F . At the same time, a new set of messages is serially loaded into the input shift registers. We note that for simpler clock tree generation, all registers in Fig. 3 are clocked by CLK_F , while CLK_S is actually implemented as a pulsed clock, which is generated using a clock-gating cell controlled by a finite state machine.

D. Decoder Complexity Analysis

In this section, we describe some hardware complexity considerations of the proposed decoder as well as the decoding latency and the throughput, which are useful for the subsequent measures for optimization.

1) *Hardware Complexity and Memory:* Each CN and VN unit of the block diagram shown in Fig. 2 consists of the processing part (logic) and the register part (memory). While the logic complexity is estimated by synthesis results, which will be discussed in Section VI, the memory complexity can easily be characterized by counting the number of required registers.

From Fig. 2, each CN unit has d_c S/P units and d_c P/S units and a set of registers for the channel LLRs. This adds

¹In this paper we target a code with $d_v = 6$ and $d_c = 32$, as opposed to [14], which targets a code with $d_v = 3$ and $d_c = 6$.

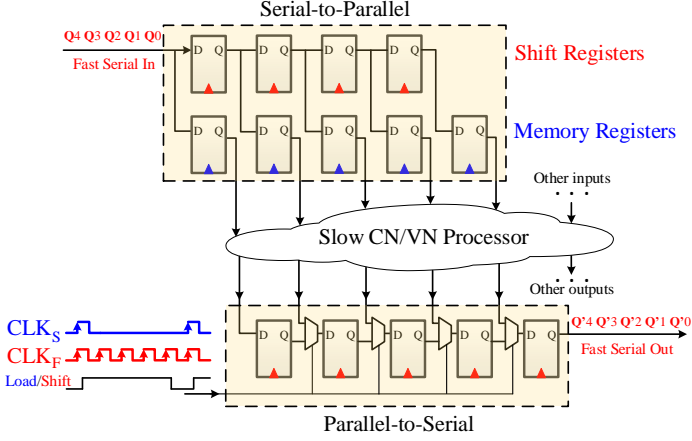


Fig. 3: The message receive and transfer mechanism by S/P and P/S shift registers for $Q_{\text{msg}} = 5$.

up to $Md_c(2Q_{\text{msg}} + Q_{\text{msg}} - 1) + 2NQ_{\text{ch}} \approx 3Md_cQ_{\text{msg}} + 2NQ_{\text{ch}}$ registers for each check node stage. Moreover, each VN unit has d_v S/P units and d_v P/S units for the message LLRs and additional registers for the channel LLRs, which corresponds to $Nd_v(2Q_{\text{msg}} + Q_{\text{msg}} - 1) + N(2Q_{\text{ch}} + Q_{\text{ch}} - 1) \approx 3Nd_vQ_{\text{msg}} + 3NQ_{\text{ch}}$ registers. The decision stage has only $Nd_v(Q_{\text{msg}} + Q_{\text{msg}} - 1) + N(Q_{\text{ch}} + Q_{\text{ch}} - 1) + N \approx 2Nd_vQ_{\text{msg}} + 2NQ_{\text{ch}}$ registers. Since the control unit registers required for each CN/VN unit are negligible compared to the registers in the data path, and with respect to the fact that $Nd_v = Md_c$ by construction of the LDPC code, the total number of registers required by the serial message-transfer decoder is

$$R_{\text{tot}} \approx (2I - 1)3N(d_vQ_{\text{msg}} + Q_{\text{ch}}) + 2N(d_vQ_{\text{msg}} + Q_{\text{ch}}), \quad (8)$$

where I is number of decoding iterations. We note that if the quantization bit-width for the channel LLRs and the message LLRs are the same, i.e., $Q_{\text{msg}} = Q_{\text{ch}} = Q$, which is often the case for MS LDPC decoders, (8) can be simplified to:

$$R_{\text{tot}} \approx N(d_v + 1)Q(6I - 1). \quad (9)$$

From (9), one can easily see that the quantization bit-width linearly increases the memory requirement for the proposed architecture.

2) *Decoding Latency*: Since each stage has a delay of two CLK_S cycles and there are two stages for each decoding iteration, the decoder latency is $4I \text{ CLK}_S$ cycles or, equivalently, $4IQ_{\text{max}} \text{ CLK}_F$ cycles, where $Q_{\text{max}} = \max(Q_{\text{msg}}, Q_{\text{ch}})$.

3) *Decoding Throughput*: In the proposed unrolled architecture, one decoder codeword is output in each CLK_S cycle. Therefore, the decoder throughput is

$$T = Nf_{\text{max}}, \quad (10)$$

where f_{max} is the maximum frequency for CLK_S . For the proposed decoder

$$f_{\text{max}} = \left\{ \max((Q_{\text{max}}T_{\text{CP,route}}), (T_{\text{CP,VN}}), (T_{\text{CP,CN}})) \right\}^{-1}, \quad (11)$$

where $T_{\text{CP,VN}}$ and $T_{\text{CP,CN}}$ are length of the critical paths of the CN unit and the VN unit, respectively, and $T_{\text{CP,route}}$ is

the critical path of the (serial) routing between the decoding stages. Thus, the decoder throughput will be limited by the routing, if the VN/CN delay is smaller than Q_{max} times the routing delay. Hence, on one hand, the serial message-transfer decoder alleviates the routing problem by reducing the required number of wires, but on the other hand, the decoder throughput for large quantization bit-widths may be affected, as the serial message-transfer delay will become the limiting factor.

IV. FINITE-ALPHABET SERIAL MESSAGE-TRANSFER LDPC DECODER

Even though the serial message-transfer architecture alleviates the routing congestion of an unrolled full-parallel LDPC decoder, it has a negative impact on both throughput and hardware complexity (more specifically memory requirement which was determined by the number of message bits), as discussed in the previous section. In our previous research [5], we have shown that finite-alphabet decoders have the potential to reduce the required number of bits while maintaining the same error rate performance. In this section, we will review the basic idea and our design method for this new type of decoders and then show how the technique can be used in order to increase the throughput and reduce the area of the serial message-transfer architecture.

A. Mutual Information Based Finite-Alphabet Decoder

In order to design a finite-alphabet LDPC decoder, we employ the approach of [5]–[7]. In this approach, custom finite-alphabet update rules that can be implemented as look-up tables (LUTs) are designed in a way that maximizes the mutual information between the VN output messages and the corresponding codeword bit.

Unfortunately, the complexity of the LUTs grows, in general, exponentially with the node degree. While [7] used monotonic LUTs for both VNs and CNs, the work of [5], [6] solved the problem for the (generally more complex) CNs by using the standard MS update rule $\Phi_c(L, \bar{\mu}) = \Phi_c^{\text{MS}}(L, \bar{\mu})$ for them. However, $\Phi_v(L_n, \bar{\mu}_{\mathcal{N}(n) \setminus m \rightarrow n})$ and $\Phi_d(L_n, \bar{\mu}_{\mathcal{N}(n) \rightarrow n})$ are designed as LUTs that preserve the ordering of the output messages in a sign-magnitude interpretation of their binary labels. The CNs can then be implemented efficiently using standard arithmetic circuits. We use this approach in the following.

Additionally, to reduce the complexity of the VN LUTs for our VN degree, we establish the proposal of [5] and [6] to use a tree of small LUTs as a more general form of the linked list of LUTs, as proposed in [7]. We note that the tree based approach requires more LUTs than the linked list approach, but it is also more parallelizable, which reduces the longest path delay. Moreover, both approaches reduce the complexity of the LUT-based VN from exponential to linear in the VN degree, with only a small negative impact on the error-correcting performance of the decoder.




Fig. 4: Frame error rate of the IEEE 802.3an LDPC code under floating-point MS decoding, fixed-point MS decoding with different bit-widths, and LUT based decoding, all with $I = 5$ decoding iterations.

B. Error-Correcting Performance and Bit-Width Reduction

In Fig. 4, we compare the performance of the IEEE 802.3an LDPC code under floating-point MS decoding, fixed-point MS decoding (with $Q_{ch} = Q_{msg} \in \{4, 5\}$), and LUT-based decoding (with $Q_{ch} = 4$ and $Q_{msg} = 3$) when performing $I = 5$ decoding iterations with a flooding schedule. We observe that the fixed-point decoder with $Q_{ch} = Q_{msg} = 5$ has almost the same performance as the floating-point decoder, while the fixed-point decoder with $Q_{ch} = Q_{msg} = 4$ shows a significant loss with respect to the floating-point implementation. Thus, a standard MS decoder would need to use at least $Q_{ch} = Q_{msg} = 5$ quantization bits. The LUT-based decoder, however, can match the performance of the floating-point decoder with only $Q_{ch} = 4$ channel quantization bits and $Q_{msg} = 3$ message quantization bits.²

C. LUT-Based Decoder Hardware Architecture

The LUT-based serial message-transfer decoder hardware architecture is very similar to the MS decoder architecture, described in Section III. However, the LUT-based decoder can take advantage of the significantly fewer message bits that need to be transferred from one decoding stage to the next. This reduction reduces the number of CLK_F cycles per iteration, which in turn increases the throughput of the decoder according to (11) provided that the CN/VN logic is sufficiently fast. Moreover, the size of the buffers needed for the S/P and P/S conversions is also reduced significantly, which directly reduces the memory complexity of the decoder (see (9)).

On the negative side, we remark that the VN units for each variable node stage (decoder iteration) of the LUT-based decoder are different, which slightly complicates the hierarchical physical implementation as we will see later. Furthermore, since $Q_{ch} > Q_{msg}$, we now need to transfer the channel LLRs with multiple (two) bits per cycle to avoid the need to artificially limit the number of CLK_F cycles per iteration to Q_{ch} rather than to the smaller Q_{msg} . To reflect this

²We note that reducing Q_{ch} further results in a non-negligible loss with respect to the floating-point decoder.

modification, we redefine (11) as $Q_{max} = \max(Q_{msg}, \lceil \frac{Q_{ch}}{2} \rceil)$. While this partially parallel transfer of the channel LLRs impacts routing congestion, we note that the overhead is negligible since the number of channel LLRs is small compared to the number of messages.

V. IMPLEMENTATION

Despite the use of a serial message-transfer, the physical implementation of the decoders, proposed in the previous sections, requires special scrutiny since the number of global wires is still significant and the area is large. Therefore, in this section, we propose a *pseudo-hierarchical* design methodology to implement the serial message-transfer architecture.

A. Physical Design Flow

Due to the large number of identical blocks in the decoder architecture, a *bottom-up* flow is expected to provide the best results. The CN, VN, and DN units are first placed and routed individually to build hard macros, and their timing and physical information are extracted. These macros are then instantiated as large cells in the decoder top level. We propose to treat the macros as *custom standard-cells* with identical height to be able to perform the placement using the standard-cell placement, rather than the less capable macro placement of the backend tool, since in our case the number of hard macro instances is extremely large and the interconnect pattern is complex and highly irregular. We also note that for the LUT-based decoder there are different macros for each variable node stage, as discussed in Section IV, while for the MS decoder an identical VN macro is used for all of the variable node stages.

Fig. 5 illustrates the proposed physical floorplan for the decoders with the unrolled architecture. In this floorplan, the CN and VN macros within each stage are constrained to be placed into dedicated regions. This measure enforces the high-level structure of the pipeline, but leaves freedom to the placement tool to choose the location for each macro in the stage to minimize routing congestion between stages. Furthermore, the CN and VN macros are placed in dedicated rows while the area between these rows is left for repeaters and for the register standard-cells for the channel LLRs in the check node stages. We note that the proposed floorplan provides the flexibility to exploit the automated algorithm for both custom and conventional standard-cells to optimize their placement in order to alleviate the still significant routing congestion.

B. Timing and Area Optimization Flow

Although the synthesis results can give an approximate evaluation for timing and area of the physical implementation, several iterations with different constraints are required to reach an optimal layout. To this end, we propose the methodology illustrated in Fig. 6 to effectively implement the serial message-transfer architecture. The main idea behind this methodology is that three main factors directly contribute to the decoder throughput and also indirectly to the decoder area, as discussed in Section III and specifically summarized in (11). Our goal is to maximize the throughput at a minimum area.




Fig. 5: The physical floorplan for serial message-transfer architecture, (a) dedicated regions for each decoder stage and (b) rows structure for macros (custom large standard-cells) and standard-cells placement.

We define the timing constraint applied to CLK_S as T_{CSTR,CLK_S} , and the timing constraint applied to CLK_F as T_{CSTR,CLK_F} . The first step is to place and route the CN/VN macros based on T_{CSTR,CLK_S} . This step is followed by the implementation of the decoder using T_{CSTR,CLK_F} . We note that the initial constraints for the backend are thereby extracted from synthesis timing results. The fully placed and routed design can give an accurate routing delay (maximum achievable CLK_F), which will be used to update T_{CSTR,CLK_F} and then T_{CSTR,CLK_S} according to (11). The updated T_{CSTR,CLK_S} will be used to re-implement the CN and VN macros within the minimum achievable area.

We note that for a strong LDPC code with a large area and long routing delay (such as the one of the IEEE 802.3an standard), the first implementation starts with $T_{CSTR,CLK_S} < Q_{max}T_{CSTR,CLK_F}$. After obtaining a realistic value for CLK_F (and hence for T_{CSTR,CLK_F}), the T_{CSTR,CLK_S} will be updated to a larger value for the next iteration. Consequently, the CN and VN macro area and thus the decoder area will shrink in the second iteration, which results in larger achievable CLK_F and hence smaller T_{CSTR,CLK_F} and T_{CSTR,CLK_S} . The feedback loop will reach the optimum point after a few iterations.

VI. RESULTS AND DISCUSSIONS

To study the impact of the serial message-transfer architecture and the finite-alphabet decoding scheme, we have implemented the proposed architecture by employing the methodology explained in Section V and analyzed the results for both MS and LUT-based decoding. We used the parity check matrix of the LDPC code defined in the IEEE 802.3an standard [15], i.e. a (2048, 1723) LDPC code of $R = 0.8413$ with $d_v = 6$ and $d_c = 32$. We used $I = 5$ for both decoders

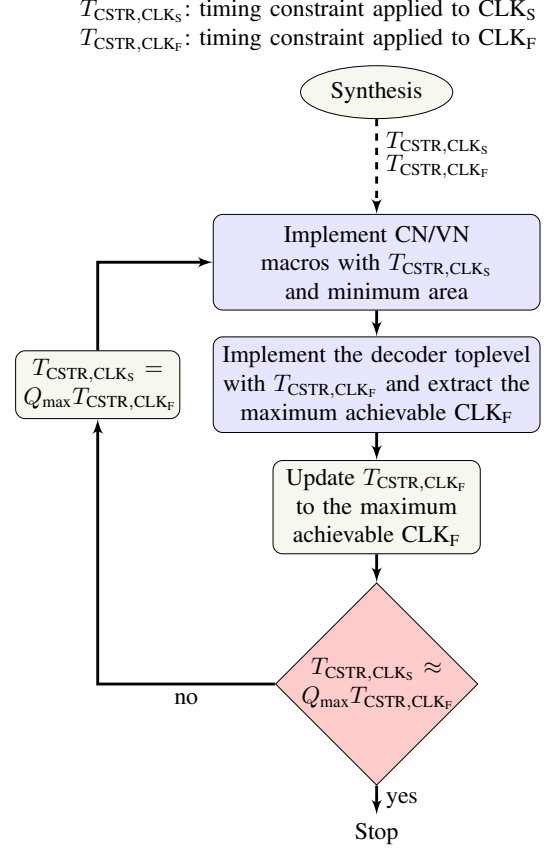


Fig. 6: The proposed algorithm to optimized timing and area for the serial message-transfer architecture.




Fig. 7: Layouts for (a) MS decoder and (b) LUT-based decoder.

and $Q_{msg} = Q_{ch} = 5$ for the MS decoder and $Q_{msg} = 3$ and $Q_{ch} = 4$ for the LUT-based decoder to achieve the same error-correction performance, as described in Sections III and IV. The decoders were synthesized from a VHDL description using Synopsys Design Compiler and placed and routed using Cadence Encounter Digital Implementation. The layouts are shown in Fig. 7. The results are reported for a 28 nm FD-SOI library under typical operating conditions ($V_{DD} = 1$ V, $T = 25^\circ$ C).

TABLE I: Critical path delays (ns) for MS and LUT-based decoder

Path	MS decoder	LUT-based decoder
CN	2.38	1.42
VN	0.96	1.24
Interconnect	1.51	1.16




Fig. 8: Detailed area results for the LUT-based decoder with total area of 16.2 mm^2 , and the MS decoder with total area of 23.3 mm^2 .

A. Delay Analysis

In the serial message-transfer architecture the critical path, and hence, the maximum decoding frequency are defined by (11). To investigate the impact of serially transferring the messages on the decoder throughput, we consider the delay of the following register-to-register path groups for both the MS and LUT-based decoder.

1) *CN Path*: The CN path is the path from the S/P memory registers to the P/S shift registers within the CN unit. For both decoders, this path is essentially comprised of the logic cells for a sorter tree with a depth of 4.

2) *VN Path*: The VN path is the path from S/P memory registers to P/S shift registers within the VN unit. This path is dominated by an adder tree for the MS decoder and an LUT tree for the LUT-based decoder.

3) *Interconnect Path*: The interconnect path comprises mainly the routing wires (and buffers) that connect the CN/VN unit S/P shift registers to the VN/CN unit P/S shift registers.

Table I summarizes the critical path delays of the CN/VN and the interconnect path groups. Together with (11), the values in the table dictate the maximum achievable frequency for CLK_S and CLK_F , respectively, for both of the decoders with the proposed serial message-transfer architecture. We note that the critical paths are reported after the timing and area constraints for the CN/VN macros and for the decoder toplevel are jointly optimized according to the flow shown in Fig. 6. From Table I, it can be seen that the CN processor

TABLE II: Detailed area (μm^2) for CN/VN unit

Component	MS decoder	LUT-based decoder
CN unit logic	1578	485
CN unit register	1695	971
CN macro	3607	1510
VN unit logic	315	403
VN unit register	381	235
VN macro	755	646

TABLE III: Power and energy efficiency comparison for MS and LUT-based decoder

	MS decoder	LUT-based decoder
Total power (mW)	12248	10257
Leakage power (mW)	7.44	5.27
Energy efficiency (pJ/bit)	45.2	22.7

in LUT-based decoder is *less* complex than the one in the MS decoder, while the the LUT-based VN processor is *more* complex compared to the one that is based on adders in the MS decoder. According to (11), we observe that in both decoders the message transfer limits the throughput to a CLK_S period of $5 \times 1.51 \text{ ns} = 7.55 \text{ ns}$ and $3 \times 1.16 \text{ ns} = 3.48 \text{ ns}$ for the MS and the LUT-based decoder, respectively, where 1.51 ns and 1.16 ns are the corresponding minimum CLK_F periods. Consequently, in our flow, the VN and CN units end up as optimized for minimum area only with relaxed and easy to meet timing constraints.

B. Area Analysis

Fig. 8 illustrates the area distribution among their components after layout. The area utilization is approximately 67% for both decoders. While almost 62% of the layout is filled with CN/VN macros and registers, clock tree and routing buffers occupy around 5%. Furthermore, we see a 44% difference in total area between the decoders due to the fact that the absolute area for CN and VN macros is 14.12 mm^2 in the MS decoder, as opposed to only 9.56 mm^2 in the LUT-based decoder.

To understand this fact, we list the area of each CN/VN macro in Table II³. According to this table, the finite-alphabet message passing algorithm leads to significantly smaller CN processors because of two important factors: first, the bit-width reduction of the messages directly affects the data-path area, and second, the quantized messages in the LUT-based decoder are processed directly in the sorter tree of the CNs without the need to compute their absolute values. However, VN processors are less area-efficient in the LUT-based decoder in comparison with the ones of the MS decoder. This is caused by the fact that the LUT-based computations are, in general, less area-efficient than the conventional arithmetic based update

³Logic and register areas are obtained from synthesis, and macro areas are the final post-layout results.

TABLE IV: Implementation results for MS and LUT-based decoder and comparison with other work

	MS decoder	LUT-based decoder	[14]	[8]	[19]
Process technology	28 nm FD-SOI		65 nm CMOS	65 nm CMOS low-power	65 nm CMOS
Supply voltage (V)	1.0		1.2	1.2	1.3
LDPC code	(2048, 1723)		(672, 546)	(2048, 1723)	(2048, 1723)
Node degree (d_v, d_c)	(6, 32)		(3, 6)	(6, 32)	(6, 32)
Algorithm	min-sum	finite-alphabet	min-sum	min-sum	split-row
Quantization bits	5	3	4	4	5
Architecture	unrolled full-parallel		unrolled full-parallel	partial-parallel	full-parallel
Core area (mm ²)	23.3	16.2	12.9	5.05	4.84
Area utilization (%)	66.4	65.9	76	84.5	97
Max. frequency (MHz)	662	862	257	700	195
Latency (ns)	151	69.6	105	137	56.4
Throughput (Gbps)	271	594	160.8	13.3*	36.3 [†]
Power @ f_{\max} (mW)	12248	10257	5360	2800	1359
Area eff. (Gbps/mm ²)	11.6	36.6	12.5	2.63	7.5
Energy per bit (pJ/bit)	45.2	22.7	33.3	210.5*	37.4 [†]
Scaled area eff. [‡] (Gbps/mm ²)	11.6	36.6	156.4	32.9	93.8
Scaled energy per bit [§] (pJ/bit)	45.2	22.7	23.1	146.2	22.1

*Throughput and energy per bit @ $I_{\max} = 8$

[†]Throughput and energy per bit @ $I_{\max} = 11$

[‡]Scaling is done by S^3 where S is relative dimension to 28

[§]Scaling is done by $1/U^2$ where U is relative voltage to 1.0

rules. Thus, the logic area of the VN in the LUT-based decoder is larger, even though their input/output bit-width is smaller. Another contributing factor in the Table II is the register area, which is defined by the number of S/P and P/S registers. For those, the 40% reduction of bit-width in the LUT-based decoder is directly noticeable in the register area for both CN and VN units. Altogether, the CN and VN macros in the LUT-based decoder are 58% and 14% smaller, respectively, compared to those of the MS decoder.

C. Power Analysis

The energy which is consumed by each decoder is proportional to the capacitance, which in turn is related to the decoder area. Also, the number of required CLK_F cycles for the serial message-transfer to decode one codeword, which is inversely proportional to the decoding throughput at a constant frequency, directly contributes to the consumed energy for each decoded bit. Therefore, we analyze both the average power and the energy efficiency of the decoders using post-layout vector-based power analysis.⁴ The results are reported in Table III. We note that the average power is calculated at a constant CLK_S frequency, here $f_F = \min(f_{\max, \text{MS}}, f_{\max, \text{LUT}}) = 662$ MHz, for both decoders. According to Table III, the total power consumption of the LUT-based decoder is 16.2% smaller than that of the MS decoder. Furthermore, by considering the fact that the LUT-based decoder has 66.7% higher throughput than the MS decoder at a similar CLK_F frequency, the energy efficiency of the LUT-based decoder is almost 2 times better in comparison with the MS decoder.

⁴We first extract the parasitic information of hard macros and toplevel from the placement and routing tool and then read and link them using a power computation tool to make the complete parasitic information.

D. Summary and Final Comparison

The final post-layout results for our MS and LUT-based decoders and also for some other recently implemented decoders are summarized in Table IV. Our LUT-based decoder runs at a maximum CLK_F frequency of $f_F = 862$ MHz and delivers 0.594 Tbps, while it occupies 16.2 mm² area and dissipates 22.7 pJ/bit. Compared to the MS decoder, the LUT-based decoder is 1.4× smaller, 2.2× faster, and thus 3.1× more area efficient. It also has 16.2% lower power dissipation and 2× better energy efficiency, when the decoding throughput is taken into account.

The work in [14] is the only other unrolled full-parallel decoder in literature, but it is designed for the IEEE 802.11ad [22] code, which has a shorter block length and smaller node degrees ($d_v = 3$ and $d_c = 6$ as opposed to $d_v = 6$ and $d_c = 32$ for the code used in the design reported in this paper). The work of [8] and [19] are for the same IEEE 802.3an code considered in this paper, but with partial-parallel and full-parallel architectures, respectively. Comparing to [8] and [19], the proposed LUT-based decoder has more than an order of magnitude higher throughput with almost the same energy efficient as the decoder in [19]. The area efficiency of the proposed unrolled full-parallel architecture, however, is inferior to the one of the decoder in [19] with full-parallel architecture due to the repeated routing overhead between the decoder stages in our design.

VII. CONCLUSION

An ultra high throughput LDPC decoder with a serial message-transfer architecture and based on non-uniform quantization of messages was proposed to achieve the highest decoding throughput in literature. The proposed decoder architecture is an unrolled full-parallel architecture with serialized

messages for CN/VN units, which was enabled by employing S/P and P/S shift registers at the inputs and outputs of each unit. The proposed quantized message passing algorithm replaces conventional MS, resulting in 40% reduction in message bit-width without any performance penalty. This algorithm was implemented by using generic LUTs instead of adders for VNs while the CNs remained unchanged compared to MS decoding. Placement and routing results in 28 nm FD-SOI show that the LUT-based serial message-transfer decoder delivers 0.594 Tbps throughput and is 3.3 times more area efficient and 2 times more energy efficient in comparison with the MS decoder with serial message-transfer architecture.

ACKNOWLEDGMENT

This work was supported by the Swiss National Science Foundation (SNSF) under the project number 200021-153640.

REFERENCES

- [1] D. J. MacKay, "Good error-correcting codes based on very sparse matrices," *IEEE Trans. Inf. Theory*, vol. 45, no. 2, pp. 399–431, 1999.
- [2] M. P. Fossorier, M. Mihaljević, and H. Imai, "Reduced complexity iterative decoding of low-density parity check codes based on belief propagation," *IEEE Trans. Commun.*, vol. 47, no. 5, pp. 673–680, 1999.
- [3] S. Planjery, D. Declercq, L. Danjean, and B. Vasic, "Finite alphabet iterative decoders – Part I: Decoding beyond belief propagation on the binary symmetric channel," *IEEE Trans. Commun.*, Oct. 2013.
- [4] B. Kurkoski, K. Yamaguchi, and K. Kobayashi, "Noise thresholds for discrete LDPC decoding mappings," in *IEEE Global Telecommunications Conf. (GLOBECOM)*, Nov. 2008.
- [5] A. Balatsoukas-Stimming, M. Meidlinger, R. Ghanaatian, G. Matz, and A. Burg, "A fully-unrolled LDPC decoder based on quantized message passing," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct 2015, pp. 1–6.
- [6] M. Meidlinger, A. Balatsoukas-Stimming, A. Burg, and G. Matz, "Quantized message passing for LDPC codes," in *Asilomar Conference on Signals, Systems and Computers*, Nov 2015, pp. 1606–1610.
- [7] F. J. C. Romero and B. M. Kurkoski, "Decoding LDPC codes with mutual information-maximizing lookup tables," in *IEEE International Symposium on Information Theory (ISIT)*, Jun. 2015, pp. 426–430.
- [8] Z. Zhang, V. Anantharam, M. J. Wainwright, and B. Nikolić, "An efficient 10GBASE-T Ethernet LDPC decoder design with low error floors," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 843–855, 2010.
- [9] A. Cevrero, Y. Leblebici, P. Jenne, and A. Burg, "A 5.35 mm² 10GBASE-T Ethernet LDPC decoder chip in 90 nm CMOS," in *IEEE Asian Solid State Circuits Conference (A-SSCC)*, IEEE, 2010, pp. 1–4.
- [10] C. Roth, P. Meinerzhagen, C. Studer, and A. Burg, "A 15.8 pJ/bit/iter quasi-cyclic LDPC decoder for IEEE 802.11n in 90 nm CMOS," in *IEEE Asian Solid State Circuits Conference (A-SSCC)*, IEEE, 2010, pp. 1–4.
- [11] T.-C. Kuo and A. N. Willson Jr, "A flexible decoder IC for WiMAX QC-LDPC codes," in *IEEE Custom Integrated Circuits Conference (CICC)*, IEEE, 2008, pp. 527–530.
- [12] A. J. Blanksby and C. J. Howland, "A 690-mW 1-Gb/s 1024-b, rate-1/2 low-density parity-check code decoder," *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 404–412, 2002.
- [13] N. Onizawa, T. Hanyu, and V. C. Gaudet, "Design of high-throughput fully parallel LDPC decoders based on wire partitioning," *IEEE Trans. VLSI Syst.*, vol. 18, no. 3, pp. 482–489, 2010.
- [14] P. Schlafer, N. Wehn, M. Alles, and T. Lehnigk-Emden, "A new dimension of parallelism in ultra high throughput LDPC decoding," in *IEEE Workshop on Signal Processing Systems (SiPS)*, IEEE, 2013, pp. 153–158.
- [15] "IEEE Standard for Information Technology – Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks – Specific Requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," IEEE Std. 802.3an, Sep. 2006.
- [16] H. Kaeslin, *Digital integrated circuit design: from VLSI architectures to CMOS fabrication*. Cambridge University Press, 2008.
- [17] A. Darabiha, A. C. Carusone, and F. R. Kschischang, "A 3.3-Gbps bit-serial block-interlaced min-sum LDPC decoder in 0.13- μ m CMOS," in *IEEE Custom Integrated Circuits Conference (CICC)*, IEEE, 2007, pp. 459–462.
- [18] A. Darabiha, A. C. Carusone, and R. Kschischang, "Power reduction techniques for LDPC decoders," *IEEE J. Solid-State Circuits*, vol. 43, no. 8, pp. 1835–1845, 2008.
- [19] T. Mohsenin, D. N. Truong, and B. M. Baas, "A low-complexity message-passing algorithm for reduced routing congestion in LDPC decoders," *IEEE Trans. Circuits Syst. I*, vol. 57, no. 5, pp. 1048–1061, 2010.
- [20] N. Mobini, A. H. Banihashemi, and S. Hemati, "A differential binary message-passing LDPC decoder," *IEEE Trans. Commun.*, vol. 57, no. 9, pp. 2518–2523, 2009.
- [21] K. Cushon, S. Hemati, C. Leroux, S. Mannor, and W. J. Gross, "High-throughput energy-efficient ldpc decoders using differential binary message passing," *IEEE Trans. Signal Process.*, vol. 62, no. 3, pp. 619–631, 2014.
- [22] "Iso/iec/ieee international standard for information technology–telecommunications and information exchange between systems–local and metropolitan area networks–specific requirements-part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 3: Enhancements for very high throughput in the 60 ghz band (adoption of ieee std 802.11ad-2012)," *ISO/IEC/IEEE 8802-11:2012/Amd.3:2014(E)*, pp. 1–634, March 2014.



Reza Ghanaatian was born in Jahrom, Iran, in 1988. He received the M.Sc. degree in digital systems from the Department of Electrical Engineering, Sharif University of Technology (SUT), Tehran, Iran, in 2012. He was with Advanced Integrated Circuit Design Laboratory (AICDL), at SUT from 2011 to 2013, working on field-programmable gate array based systems for wireless and optical communication applications.

In 2014, Mr. Ghanaatian joined Telecommunications Circuits Laboratory (TCL) at EPFL, Lausanne, Switzerland, working towards his Ph.D. His current research interests include VLSI circuits for signal processing and communications as well as approximate computing techniques for energy efficient system design.



Alexios Balatsoukas-Stimming received the Diploma and M.Sc. degrees in Electronics and Computer Engineering from the Technical University of Crete, Greece, in 2010 and 2012, respectively. His M.Sc. studies were supported by a scholarship from the Alexander S. Onassis foundation. He received the Ph.D degree in Computer and Communications Sciences from EPFL, Switzerland, where he performed his research at the Telecommunications Circuits Laboratory.

He serves as reviewer for several IEEE Journals and Conferences, and has been recognized as an Exemplary Reviewer by the IEEE Wireless Communications Letters in 2013 and 2014. His current research interests include VLSI circuits for signal processing and communications, as well as error correction coding theory and practice.



Christoph Müller was born in Leipzig, Germany, in 1986. He received the bachelors degree in information technology from the Schmalkalden University of Applied Sciences, Schmalkalden, Germany, in 2011, and the masters degree in System-on-Chip from Lund University, Lund, Sweden, in 2013. He has been working as Project/Research Assistant at Lund University and at Technical University of Denmark, Kongens Lyngby, Denmark. Currently he is pursuing his PhD degree at EPFL, Lausanne, Switzerland with a research focus on digital implementation and an emphasis on low power, variation mitigation and design methodology.



Michael Meidlinger was born in Austria in 1989. He received his Bsc. and Msc. degrees from Technische Universität (TU) Wien, Vienna, Austria in 2011 and 2013 respectively, both with distinction. From 2011 to 2013, Michael has been working in the field of mobile communication research and helped to develop the Vienna LTE-A Simulators. Since 2013, Michael is part of the Communication Theory group at TU Wien, working towards his Ph.D. His current research interests include quantizer design for telecommunication receivers as well as error

correction coding and superposition modulation techniques.



Gerald Matz (S'95M'01SM'07) received the Dipl.-Ing. degree (1994) and the Dr. techn. degree (2000) in Electrical Engineering and the Habilitation degree (2004) for Communication Systems from Vienna University of Technology, Austria. He is currently an Associate Professor with the Institute of Telecommunications, Vienna University of Technology. He has held visiting positions with the Laboratoire des Signaux et Systèmes, Ecole Supérieure d'Électronique, France; with the Communication Theory Lab at ETH Zurich, Switzerland; and with Ecole Nationale Supérieure d'Électronique, d'Informatique et d'Hydraulique de Toulouse, France. He has published some 170 scientific articles in international journals, conference proceedings, and edited books. He is co-editor of the book *Wireless Communications over Rapidly Time-Varying Channels* (New York: Academic, 2011). His research interests include wireless communications, sensor networks, statistical signal processing, and information theory. Prof. Matz was on the IEEE SPS Technical Committee on Signal Processing for Communications and Networking, and currently serves on the IEEE SPS Technical Committee on Signal Processing Theory and Methods. He is Associate Editor of the IEEE TRANSACTIONS ON INFORMATION THEORY and was on the Editorial Board of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the EURASIP Journal Signal Processing, and the IEEE SIGNAL PROCESSING LETTERS. He was Lead Guest Editor of the Special Issue on Managing Complexity in Multiuser MIMO Systems in the IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING, Technical Program Co-Chair of EUSIPCO 2004, and member of the Technical Program Committee of numerous international conferences. In 2006, he received the Kardinal Innitzer Most Promising Young Investigator Award. He is a Senior Member of the IEEE.

He is a Senior Member of the IEEE.



Adam Teman received the Ph.D., M.Sc., and B.Sc. degrees in Electrical Engineering from Ben-Gurion University (BGU), Be'er Sheva, Israel in 2006, 2011, and 2014, respectively. He worked as a Design Engineer at Marvell Semiconductors from 2006 to 2007, with an emphasis on Physical Implementation. Dr. Teman's research interests include low-voltage digital design, energy efficient SRAM, NVM, and eDRAM memory arrays, low power CMOS image sensors, low power design techniques for digital and analog VLSI chips, significance-driven approximate

computing, and process tolerant design techniques. He has authored more than 40 scientific papers and 4 patent applications and is an associate editor at the Microelectronics Journal and a technical committee member of several IEEE conferences. In 2010–2012, Dr. Teman was honored with the Electrical Engineering Department's Teaching Excellence recognition at BGU, and in 2011, he was awarded with BGU's Outstanding Project award. Dr. Teman received the Yizhak Ben-Yaakov HaCohen Prize in 2010, the BGU Rector's Prize for Outstanding Academic Achievement in 2012, the Wolf Foundation Scholarship for excellence of 2012 and the Intel Prize for Ph.D. students in 2013. His doctoral studies were conducted under a Kreitman Foundation Fellowship. Dr. Teman was a post-doctoral researcher at the Telecommunications Circuits Lab (TCL) at the École Polytechnique Fédérale de Lausanne (EPFL), Switzerland under a Swiss Government Excellence Scholarship from 2014–2015. In October 2015, Dr. Teman joined the faculty of engineering at Bar-Ilan University, Ramat Gan, Israel in 2015, where he is currently a tenure track researcher in the department of electrical engineering and a leading member of the Emerging Nanoscaled Integrated Circuits and Systems (EnICS) Labs.



Andreas Burg (S'97-M'05) was born in Munich, Germany, in 1975. He received his Dipl.-Ing. degree from the Swiss Federal Institute of Technology (ETH) Zurich, Zurich, Switzerland, in 2000, and the Dr. sc. techn. degree from the Integrated Systems Laboratory of ETH Zurich, in 2006. In 1998, he worked at Siemens Semiconductors, San Jose, CA. During his doctoral studies, he worked at Bell Labs Wireless Research for a total of one year. From 2006 to 2007, he was a postdoctoral researcher at the Integrated Systems Laboratory and at the

Communication Theory Group of the ETH Zurich. In 2007 he co-founded Celestris, an ETH-spinoff in the field of MIMO wireless communication, where he was responsible for the ASIC development as Director for VLSI. In January 2009, he joined ETH Zurich as SNF Assistant Professor and as head of the Signal Processing Circuits and Systems group at the Integrated Systems Laboratory. Since January 2011, he has been a Tenure Track Assistant Professor at the École Polytechnique Fédérale de Lausanne (EPFL) where he is leading the Telecommunications Circuits Laboratory.

In 2000, Mr. Burg received the Willi Studer Award and the ETH Medal for his diploma and his diploma thesis, respectively. Mr. Burg was also awarded an ETH Medal for his Ph.D. dissertation in 2006. In 2008, he received a 4-years grant from the Swiss National Science Foundation (SNF) for an SNF Assistant Professorship. With his students he received the best paper award from the EURASIP Journal on Image and Video Processing in 2013 and best demo/paper awards at ISCAS 2013, ICECS 2013, and at ACSSC 2007.

He has served on the TPC of various conferences on signal processing, communications, and VLSI. He was a TPC co-chair for VLSI-SoC 2012 and the TCP co-chair for ESSCIRC 2016 and SiPS 2017. He served as an Editor for the IEEE Transaction of Circuits and Systems in 2013 and is on the Editorial board of the Springer Microelectronics Journal and the MDPI Journal on Low Power Electronics and its Applications. He is also a member of the EURASIP SAT SPCN and of the IEEE TC-DISPS.